JB 90-02

Paper Entitled:

"Bayesian Sequential Hypothesis Testing",

(with D. C. MacEnany),

Proceedings of:

11th IFAC World Congress
August 1990, Tallian
USSR

Vol. 3, pp. 227-232

BAYESIAN SEQUENTIAL HYPOTHESIS TESTING

John S. Baras

David C. MacEnany

Electrical Engineering Department and Systems Research Center, University of Maryland, College Park, Maryland 20742, USA

Abstract. We consider the Bayesian sequential detection problem for general observation processes with a continuous time parameter. Primary emphasis is placed on first exit policies and their generic optimality. A new geometric formulation and solution to the existence and uniqueness of the optimal first exit policy will be given as well as an explicit constructive algorithm for its computation. We will apply these results to the problem of Bayesian sequential detection for diffusion-type observations although our methods and results generalize to problems of optimal stopping and decision for partially observed Markov chains where the observation processes are of the diffusion or point process type. These results will appear elsewhere.

Keywords. Bayesian statistics; boundary-value problems; optimal control; optimal stopping policies; signal detection.

PROBABILISTIC FRAMEWORK AND MAIN RESULTS

Let's begin with a precise statement of the Bayesian formulation for the sequential detection problem of two simple hypotheses. Let $(\Theta, \mathcal{P}(\Theta), P)$ be a probability triple carrying a binary-valued random variable, H_π, where $\Theta = \{0,1\}$, $\mathcal{P}(\Theta)$ denotes the power set of Θ, and where P is a probability measure assigning mass to H_π according to,

$$P\{\theta \in \Theta : H_\pi(\theta) = 1\} = \pi \tag{1.1}$$

for some $\pi \in [0,1]$. Employing an abbreviated notation we have immediately, $P\{H_\pi = 0\} = 1 - \pi$. Next, let $(\Omega_o, \mathcal{O})$ be a measurable space upon which there are defined two probability measures, P_0, P_1, which will serve to model the statistics of the observable events in $\mathcal{O}$ described by $\{\mathcal{O}_t : t \geq 0\}$, a right-continuous filtration on $\mathcal{O}$ to which we assume we have access and for which $\mathcal{O}_0 = \{\phi, \Omega_o\}$. From these probability triples construct the product triple $(\Omega, \mathcal{F}, P_\pi)$ via $\Omega = \Theta \times \Omega_o$, $\mathcal{F} = \mathcal{P}(\Theta) \otimes \mathcal{O}$, and with the probability measure P_π satisfying,

$$P_\pi\{\{H_\pi = i\} \cap O\} = P\{H_\pi = i\} P_i\{O\} \quad \forall O \in \mathcal{O};\ i = 0,1, \tag{1.2}$$

which yields,

$$P_\pi\{O\} = \pi P_1\{O\} + (1-\pi) P_0\{O\} \qquad \forall O \in \mathcal{O}. \tag{1.3}$$

Within this set-up we assume that the random variable H_π is unobservable but that one can observe some stochastic process with filtration $\{\mathcal{O}_t : t \geq 0\}$ whose statistics in the events of $\{H_\pi = 0\}, \{H_\pi = 1\} \in \Theta$ are governed by the probability measures P_0, P_1, respectively. In the task of trying to determine which event set is responsible for what is observed, we employ a two-part decision structure. First, a decision to terminate the observation of the process is made according to an $(\mathcal{O}_t, P_\pi)$-stopping time, say τ. Second, an inference as to the true value of H_π is made according an $(\mathcal{O}_\tau, P_\pi)$-binary-valued random variable, say δ. Any such pair, (τ, δ), is called an *admissible policy*. Over the set of admissible policies we define a cost function, $\rho_\pi(\tau,\delta)$, usually called *Bayes' cost*,

$$\rho_\pi(\tau,\delta) := E_\pi\left[\int_0^\tau c_s\, ds + C(H_\pi, \delta)\right] \tag{1.4}$$

where $\{c_t : t \geq 0\}$ is some nonnegative, $\mathcal{O}_t$-adapted process, and where for $c^0, c^1 > 0$ we define,

$$C(H,\delta) := \begin{cases} c^0, & \text{if } H = 1 \ \&\ \delta = 0; \\ 0, & \text{if } H = \delta; \\ c^1, & \text{if } H = 0 \ \&\ \delta = 1. \end{cases} \tag{1.5}$$

We make the following technical assumptions concerning $\{c_t : t \geq 0\}$,

(A1): $E_\pi \int_0^t c_s ds < \infty \qquad \forall t < \infty$.

(A2): $P_\pi\left\{\int_0^\infty c_s ds = \infty\right\} = 1$;

The first assumption reflects the desire that there be no *a priori* fixed minimum amount of time before which a policy must make a decision. The second assumption removes from consideration those policies which have even the slightest chance of not making a decision in a finite amount of time.

We see that $E_\pi[C(H_\pi,\delta)]$, called the average *terminal cost*, yields the P_π-average cost of an incorrect decision, while $E_\pi \int_0^\tau c_s\, ds$, called the average *running cost*, is interpreted as the P_π-average cost of not making a decision until the $(\mathcal{O}_t, P_\pi)$-stopping time τ; Bayes' cost is the sum of the two. We point out that it is without loss of generality that no cost is levied for correct decisions in 1.5.

Having gotten this far, we can now most succinctly state that our modest goal is to find an admissible policy, (τ_*, δ_*), which minimizes 1.4 over the set of all admissible policies. This leads to our first definition.

Definition 1. An admissible policy, (τ_*, δ_*), is said to be *Bayesian optimal* if,

$$\rho_\pi(\tau_*, \delta_*) = \inf_{(\tau,\delta)} \rho_\pi(\tau,\delta) \qquad \forall \pi \in [0,1], \tag{1.6}$$

where the infimum is over all admissible policies. We define $\rho(\pi) := \inf_{(\tau,\delta)} \rho_\pi(\tau_*, \delta_*)$ and call it Bayes' optimal cost. □

If we use 1.5, 1.2, and 1.3 to rewrite Bayes' cost as,

$$\begin{aligned} \rho_\pi(\tau,\delta) &= E_\pi\Big[\int_0^\tau c_s ds + c^0\, 1\{H_\pi = 1, \delta = 0\} + c^1\, 1\{H_\pi = 0, \delta = 1\}\Big] \\ &= \pi E_1\left[\int_0^\tau c_s\, ds + c^0\, 1\{\delta = 0\}\right] \\ &\quad + (1-\pi) E_0\left[\int_0^\tau c_s\, ds + c^1\, 1\{\delta = 1\}\right], \end{aligned} \tag{1.7}$$

we see that $\rho_\pi(\tau,\delta)$ is linear in π and therefore Bayes' optimal cost, $\rho(\pi)$, being the lower envelope of a family of lines is necessarily concave. From this definition it is also clear why we assume *both* c^0 and c^1 are strictly positive. For if $c^0 c^1 = 0$ then it follows easily that $\rho \equiv 0$, and the trivial policy (τ_*, δ_*) with $\tau_* \equiv 0$ $\delta_* = 1\{c^1 = 0\}$ is seen to be Bayesian optimal.

The minimization in 1.6 can be greatly simplified as follows. Define the $(\mathcal{O}_t, P_\pi)$-conditional probability of the 'hypothesized' event, $\{H_\pi = 1\}$ as,

$$\Pi_t := P_\pi\{H_\pi = 1|\mathcal{O}_{t-}\} \qquad t \geq 0. \tag{1.8}$$

Note that this definition implies $P_\pi\{\Pi_0 = \pi\} = 1$ and that the Π process is a corlol uniformly integrable $(\mathcal{O}_t, P_\pi)$-predictable martingale. Consider the following lemma, proven for the discrete-time case in (Shiryayev, 1978) and as stated below in (MacEnany, 1987).

Lemma 1 (Optimal Stopping): Define,

$$\rho_\pi(\tau) := E_\pi\left[\int_0^\tau c_s\, ds + e(\Pi_\tau)\right], \tag{1.9}$$

with,

$$e(\pi) := \min\{c^0\pi, c^1(1-\pi)\}. \tag{1.10}$$

Let $\mathcal{T}$ denote the class of $(\mathcal{O}_t, P_\pi)$-a.s. finite stopping times. Then,

$$\inf_{\tau\in\mathcal{T}} \rho_\pi(\tau) = \inf_{(\tau,\delta)} \rho_\pi(\tau,\delta), \tag{1.11}$$

where the infimum on the right is over all admissible policies.

As a result, *the search for an optimal policy can be reduced to a search for an optimal stopping time.* Note that according to our definition of ρ in Definition 1 we have shown, $\rho(\pi) = \inf_{\tau\in\mathcal{T}} \rho_\pi(\tau)$. The next step is to transform this minimization problem into one still more manageable. In trying to do this it becomes clear that a most important subclass of admissible stopping times are those which are first exit times of the Π process from an interval. This is due both to their simple specification and remarkable optimality properties.

Definition 2: The first exit time, τ_I, of Π from an interval I is a $(\mathcal{O}_t, P_\pi)$-stopping time defined as,

$$\tau_I := \inf\{t \geq 0 : \Pi_t \notin I\}, \tag{1.12}$$

where $I \subset [0,1]$, nonempty, is called the continuation interval. We will denote by $\overline{\mathcal{T}}$ the collection of such first exit times. The admissible policy given by $(\tau_I, \delta_*(\Pi_{\tau_I}))$ will be called a first exit policy; we observe that it is an admissible policy. □

Our plan of course is to simplify the minimization over $\mathcal{T}$ by replacing it with a minimization over $\overline{\mathcal{T}}$. Prerequisite to the success of this plan are conditions which guarantee that a first exit policy based on $I \subset [0,1]$ satisfies $\tau_I \in \mathcal{T}$, i.e., which guarantee that Π eventually (P_π-a.s.) exits I. To state such *escape conditions* in general will take us too far afield for our present purposes and *we will assume for the remainder of this section that the first exit times under consideration are in fact P_π-a.s. finite.* We stress that without such conditions one cannot be sure *a priori* that $\tau_I \in \mathcal{T}$. For instance, although $\tau_I \in \overline{\mathcal{T}}$ is clearly an $\mathcal{O}_t$-stopping time for any open interval $I \subset [0,1]$, we have given no guarantee that such a stopping time is P_π-a.s. finite. The interested reader can consult (Liptser and Shiryayev, 1978; MacEnany, 1987) for a full treatment of general escape conditions. For our present purposes we will defer stating escape conditions until the specifics of the applications to follow.

Still, when Π exits an interval, questions naturally arise as to its whereabouts at the time of escape. The next definition provides a means to phrase such questions.

Definition 3: Let I be a continuation interval containing $\pi \in I$ and suppose that $\tau_I \in \mathcal{T}$. Let $\overline{\Omega} = \{\omega \in \Omega : \Pi_0(\omega) = \pi\}$. We define the Π-*boundary of* I as,

$$\partial_\Pi I := \bigcup_{\omega\in\overline{\Omega}} \{\pi \notin I : \pi = \Pi_{\tau_I(\omega)}(\omega)\},$$

and the Π-*closure of* I as, $[I]_\Pi := I \bigcup \partial_\Pi I$. We point out that assumption $\pi \in I$ leads to the implication $P_\pi\{\Pi_0 \in I\} = 1$ in view of the fact that $P_\pi\{\Pi_0 = \pi\} = 1$. From this and the supposition $\tau_I \in \mathcal{T}$ it follows that $P_\pi\{\Pi_{\tau_I} \in \partial_\Pi I\} = 1$. □

If the sample paths of Π are continuous, we see that for any interval I that $\partial_\Pi I \subseteq \partial I$, where ∂I denotes the usual boundary of I, i.e., its endpoints. However this is not usually true if Π has jumps. We are now in a position to state a set of conditions which if true will imply that, $\inf_{\tau\in\mathcal{T}} \rho_\pi(\tau) = \inf_{\tau\in\overline{\mathcal{T}}} \rho_\pi(\tau)$, which is a considerable and welcome simplification.

Let $I_* \subset [0,1]$ be some nonempty interval, let r_* denote a mapping, $r_* : [0,1] \to \mathbf{R}$, and consider the following conditions on the pair (r_*, I_*):

$$(C1) \qquad E_\pi[r_*(\Pi_\tau) - r_*(\Pi_0)] = -E_\pi \int_0^\tau c_s\, ds \qquad \forall \tau \in \mathcal{T}$$

$$(C2) \qquad r_*(\pi) = e(\pi) \qquad \forall \pi \in \partial_\Pi I_*$$

$$(C3) \qquad r_*(\pi) < e(\pi) \qquad \forall \pi \notin \partial_\Pi I_*$$

We have the following theorem.

Theorem 1 (Verification): Suppose there exists a pair (r_*, I_*) satisfying ($C1$-$C3$), and suppose Π satisfies the *escape conditions*. Then,

$$\rho(\pi) = \inf_{\tau\in\overline{\mathcal{T}}} \rho_\pi(\tau). \tag{1.13}$$

Proof: From ($C2$) and ($C3$) we have,

$$E_\pi[e(\Pi_\tau) - r_*(\Pi_\tau)] \geq 0 \qquad \forall \tau \in \mathcal{T}. \tag{1.14}$$

Define $\tau_* \in \overline{\mathcal{T}}$ via $\tau_* := \tau_{I_*}$ and assume that $\pi \in [I_*]_\Pi$ then,

$$E_\pi[e(\Pi_{\tau_*}) - r_*(\Pi_{\tau_*})] = 0. \tag{1.15}$$

Moreover, the *escape conditions* imply $\tau_* \in \mathcal{T}$ and therefore 1.14 and 1.15 yield,

$$\inf_{\tau\in\mathcal{T}} E_\pi[e(\Pi_\tau) - r_*(\Pi_\tau)] = 0 = \inf_{\tau\in\overline{\mathcal{T}}} E_\pi[e(\Pi_\tau) - r_*(\Pi_\tau)]. \tag{1.16}$$

Hence using ($C1$) we obtain,

$$\begin{aligned} \rho(\pi) = \inf_{\tau\in\mathcal{T}} \rho_\pi(\tau) &= \inf_{\tau\in\mathcal{T}} E_\pi\left[\int_0^\tau c_s\, ds + e(\Pi_\tau)\right] \\ &= \inf_{\tau\in\mathcal{T}} E_\pi[r_*(\Pi_0) - r_*(\Pi_\tau) + e(\Pi_\tau)] \\ &= r_*(\pi) + \inf_{\tau\in\mathcal{T}} E_\pi[e(\Pi_\tau) - r_*(\Pi_\tau)] \\ &= r_*(\pi), \end{aligned} \tag{1.17}$$

where the second to last line follows the fact that $P_\pi\{\Pi_0 = \pi\} = 1$. Continuing from 1.17 we have,

$$\begin{aligned} \rho(\pi) = \inf_{\tau\in\mathcal{T}} \rho_\pi(\tau) &= r_*(\pi) + \inf_{\tau\in\overline{\mathcal{T}}} E_\pi[e(\Pi_\tau) - r_*(\Pi_\tau)] \\ &= \inf_{\tau\in\overline{\mathcal{T}}} E_\pi[r_*(\Pi_0) - r_*(\Pi_\tau) + e(\Pi_\tau)] \\ &= \inf_{\tau\in\overline{\mathcal{T}}} E_\pi\left[\int_0^\tau c_s\, ds + e(\Pi_\tau)\right] \\ &= \inf_{\tau\in\overline{\mathcal{T}}} \rho_\pi(\tau). \end{aligned} \tag{1.18}$$

In summary, 1.17 and 1.18 yield,

$$\rho(\pi) = \inf_{\tau\in\mathcal{T}} \rho_\pi(\tau) = \inf_{\tau\in\overline{\mathcal{T}}} \rho_\pi(\tau) = r_*(\pi) \qquad \forall \pi \in [I_*]_\Pi, \tag{1.19}$$

Combining 1.19 and ($C2$) we observe that,

$$\rho(\pi) = e(\pi) \qquad \forall \pi \in \partial_\Pi I_*. \tag{1.20}$$

In fact, 1.20 immediately implies that,

$$\rho(\pi) = e(\pi) \qquad \forall \pi \notin I_*. \tag{1.21}$$

This is true because ρ is concave and $\rho(0) = \rho(1) = 0$, while e is linear outside of I_* and $e(0) = e(1) = 0$. On the other hand suppose that $\pi \notin I_*$ then,

$$\begin{aligned} \rho_\pi(\tau_*) &= E_\pi\left[\int_0^{\tau_*} c_s\, ds + e(\Pi_{\tau_*})\right] \\ &= E_\pi[0 + e(\Pi_0)] = e(\pi) \end{aligned} \tag{1.22}$$

and this together with 1.21 yields $\rho(\pi) = \rho_\pi(\tau_*)$ for all $\pi \notin I_*$. From 1.21 and 1.22 there follows,

$$\rho(\pi) = \inf_{\tau\in\mathcal{T}} \rho_\pi(\tau) = \inf_{\tau\in\overline{\mathcal{T}}} \rho_\pi(\tau) = e(\pi) \qquad \forall \pi \notin I_*. \tag{1.23}$$

Combining 1.19 and 1.23 gives what we want, namely,

$$\rho(\pi) = \inf_{\tau\in\mathcal{T}} \rho_\pi(\tau) = \inf_{\tau\in\overline{\mathcal{T}}} \rho_\pi(\tau) \qquad \forall \pi \in [0,1]. \tag{1.24}$$

□

Corollary 1: The first exit policy with continuation interval I_* is Bayesian optimal for all priors $\pi \in [0,1]$. Moreover, Bayes' optimal risk is given by,

$$\rho(\pi) = \begin{cases} r_*(\pi) & \text{if } \pi \in I_* \\ r_*(\pi) = e(\pi) & \text{if } \pi \in \partial_\Pi I_* \\ e(\pi) & \text{if } \pi \notin I_* \end{cases} \tag{1.25}$$

Proof: Contained in the proof of the theorem is the fact that,

$$\rho(\pi) = \inf_{\tau\in\overline{\mathcal{T}}} \rho_\pi(\tau) = \rho_\pi(\tau_*), \qquad \forall \pi \in [0,1] \tag{1.26}$$

and hence the first exit policy (τ_*, δ_*) with $\tau_* = \tau_{I_*}$ and $\delta_* = \delta_*(\Pi_{\tau_*})$ is optimal. The expression 1.25 is simply an explicit restatement of 1.19 and 1.23. □

Corollary 2: The mapping r_* is concave on $[I_*]_\Pi$.

Proof: This follows immediately from 1.19 and the fact that ρ is concave. □

Thus, the theorem tells us that we need only search for an optimal exit time from amongst candidates in $\overline{\mathcal{T}}$, provided there exists a pair (r_*, I_*) satisfying $(C1$ - $C3)$, and what's more, if we can solve the problem posed by $(C1$ - $C3)$ then this yields the optimal (first exit) policy and our search is over. The attentive reader will have noticed that we made no use of the strictness of the inequality in $(C3)$. However, explicit usage of it is made in the next result which states that the pair (r_*, I_*) is *essentially unique*, i.e., if there exists another pair, say (s_*, J_*), satisfying $(C1$ - $C3)$ then,

$$\begin{cases} r_*(\pi) = s_*(\pi) & \forall \pi \in [I_*]_\Pi; \text{ and,} \\ I_* \equiv J_*. \end{cases} \tag{1.27}$$

We point out that the possible lack of uniqueness of r_* outside of I_* is irrelevant to any questions concerning Bayes' optimal cost or the optimal first exit policy. For our purposes, only 'uniqueness' according to 1.27 is 'essential'. Before giving this uniqueness result, we prove the following necessary condition on I_*.

Lemma 2: Under the conditions $(C1$ - $C3)$, the maximum value of the mapping e is contained in I_*, i.e., $\pi_e \in I_*$.

Proof: Let $I_* = (a_*, b_*)$ and suppose that $a_* \geq \pi_e$. Choose $\pi \in I_*$ and pick $\lambda \in (0,1)$ to satisfy $\pi = \lambda a_* + (1-\lambda)b_*$. From the second corollary to Theorem 1 we know that r_* is concave on I_* (a fortiori) and therefore,

$$\begin{aligned} r_*(\pi) &\geq \lambda r_*(a_*) + (1-\lambda) r_*(b_*) \\ &= \lambda e(a_*) + (1-\lambda) e(b_*), \end{aligned} \tag{1.28}$$

using condition $(C2)$. Since $\pi_e \notin I_*$, e is linear on I_* which implies,

$$e(\pi) = \lambda e(a_*) + (1-\lambda) e(b_*). \tag{1.29}$$

Combining 1.28, 1.29 and utilizing the arbitrariness of $\pi \in I_*$ we have,

$$r_*(\pi) \geq e(\pi) \qquad \forall \pi \in I_*. \tag{1.30}$$

But according to $(C3)$,

$$r_*(\pi) < e(\pi) \qquad \forall \pi \notin \partial_\Pi I_*, \tag{1.31}$$

which is a clear contradiction to 1.30. As a result, we must have $a_* < \pi_e$. By a similar argument we derive a contradiction unless $b_* > \pi_e$. Hence, we insist that,

$$a_* < \pi_e < b_* \tag{1.32}$$

and the lemma is shown. □

Theorem 2 (Essential Uniqueness): If a pair, (r_*, I_*), exists satisfying $(C1$ - $C3)$, then it is essentially unique.

Proof: Suppose there exists another pair which satisfies $(C1$ - $C3)$, say (s_*, J_*). Let $I_* = (a_*, b_*)$, and $J_* = (c_*, d_*)$; we will derive 1.36. From the previous lemma we know that $I_* \cap J_* \neq \phi$. Hence, if $a_* \neq c_*$ then either $c_* \in (a_*, b_*)$ or $a_* \in (c_*, d_*)$. Both possibilities lead to a contradiction. For instance, if $c_* \in (a_*, b_*)$ then according to Corollary 1 and condition $(C3)$ there holds,

$$\rho(c_*) = r_*(c_*) < e(c_*). \tag{1.33}$$

But Corollary 1 applied to s_* implies,

$$\rho(c_*) = s_*(c_*) = e(c_*). \tag{1.34}$$

A similar contradiction is obtained if $a_* \in (c_*, d_*)$ and hence, $a_* = c_*$. By an analogous argument, $b_* \neq d_*$ is untenable and therefore $I_* \equiv J_*$ which is half of *essential uniqueness*. Now with this in mind, applying Corollary 1 again yields,

$$r_*(\pi) = \rho(\pi) = s_*(\pi) \qquad \forall \pi \in [I_*]_\Pi, \tag{1.35}$$

which is the other half. □

SEQUENTIAL DETECTION FOR DIFFUSIONS

Given the probabilistic set-up of the previous section, a particular sequential detection application begins by specifying the nature of $\{O_t : t \geq 0\}$, the family of σ-algebras generated by the observations. In this application section we assume that we observe a diffusion process, Y, whose statistics conditional upon the events $\{H_\pi = 0\}, \{H_\pi = 1\}$ are governed by the probability measures P_0 and P_1, respectively. In particular, for each $\omega = (\theta, \omega_o) \in \Omega$, the observation process is given for all $t \geq 0$ by,

$$Y_t(\omega) := \begin{cases} W_t(\omega_o), & \text{if } \theta \in \{H_\pi = 0\}; \\ \int_0^t A_s(\omega_o) ds + W_t(\omega_o), & \text{if } \theta \in \{H_\pi = 1\}, \end{cases} \tag{2.1}$$

where W is a standard Wiener process, and where A is an $\mathcal{F}_t$-progressive process satisfying,

$$E_i \int_0^t |A_s| \, ds < \infty \qquad \forall t < \infty;\ i = 0, 1. \tag{2.2}$$

Typically, A represents a memoryless, nonlinear transformation of some unobserved state process driven by white noise. As indicated above, we define $O_t := \sigma\{Y_s : 0 \leq s \leq t\}$ for all $t \geq 0$. As is usual in the Bayesian formulation of a detection problem, we assume *a priori* knowledge of the probabilities that the detector is processing noise alone or signal plus noise. Therefore, in keeping with the previous section we assume that $P\{H_\pi = 1\} = \pi$ for some $\pi \in [0, 1]$. We choose the running cost for this problem as,

$$\int_0^\tau c_s \, ds = c \int_0^\tau \hat{A}_s^2 \, ds, \quad \forall \tau \in \mathcal{T};\ c > 0. \tag{2.3}$$

Since the Π process is central to our investigation, our next task is to compute a martingale representation for it. From 2.1 it is clear that $P_0 \ll P_1$, so let

$$\Lambda_t(\omega) = E_0 \left[\frac{dP_1}{dP_0} | O_t \right] \tag{2.4}$$

denote the likelihood ratio for the problem. In the Appendix we show that,

$$\Pi_t = \frac{\pi \Lambda_t}{1 - \pi + \pi \Lambda_t}, \tag{2.5}$$

and so, applying the Itô formula,

$$\Pi_t - \Pi_0 = \int_0^t \frac{\Pi_s(1-\Pi_s)}{\Lambda_s} d\Lambda_s - \int_0^t \frac{\Pi_s^2(1-\Pi_s)}{\Lambda_s^2} d[\Lambda, \Lambda]_s, \tag{2.6}$$

where the co-quadratic variation process $[X, Y]$ is given directly by

$$[X, Y]_t = X_t Y_t - \int_0^t X_{s-} dY_s - \int_0^t Y_{s-} dX_s \qquad t \geq 0 \tag{2.7}$$

for any two real semimartingales (Wong and Hajek, 1985). The likelihood ratio for this problem is well known (Wong and Hajek, 1985) and is given by,

$$\Lambda_t = \exp\left\{ \int_0^t \hat{A}_s \, dY_s - \frac{1}{2} \int_0^t \hat{A}_s^2 ds \right\} \quad t \geq 0, \tag{2.8}$$

and satisfies,

$$\Lambda_t = 1 + \int_0^t \hat{A}_s \Lambda_s \, dY_s \qquad t \geq 0, \tag{2.9}$$

where, $\hat{A}_t := E_1[A_t | O_t]$. In the proposition in the Appendix we show,

$$E_\pi[H A_t | O_t] = \Pi_t \hat{A}_t \qquad t \geq 0, \tag{2.10}$$

from which it follows that there exists a P_π-standard Wiener martingale, say $\overline{W}$, such that Y has the stochastic differential,

$$dY_t = \Pi_t \hat{A}_t \, dt + d\overline{W}_t. \tag{2.11}$$

Combining 2.6, 2.9, and 2.11 yields,

$$\begin{aligned} \Pi_t - \Pi_0 &= \int_0^t \Pi_s(1-\Pi_s) \hat{A}_s dY_s - \int_0^t \Pi_s^2 (1-\Pi_s) \hat{A}_s^2 ds \\ &= \int_0^t \Pi_s(1-\Pi_s) \hat{A}_s d\overline{W}_s \qquad \forall t \geq 0, \end{aligned} \tag{2.12}$$

Because $\overline{W}$ is a continuous martingale it is clear from 2.11 that $\partial_\Pi I \subseteq \partial I$, for $I \subset (0,1)$ any nonempty, open proper subinterval. However, it may very well be that the Π-boundary of I is empty unless we know that Π almost surely escapes I in finite time. If the following *escape conditions* are met then Π is guaranteed to exit such an interval in finite-time,

(E1): $P_i\{\int_0^\infty \hat{A}_s^2 \, ds = \infty\} = 1 \quad i = 0, 1;$

(E2): $P_i\{\int_0^t \hat{A}_s^2 \, ds = \infty\} = 0 \quad i = 0, 1;\ \forall t < \infty.$

We do not prove their validity here but instead refer the reader to (MacEnany 1987). Escape condition (E1) has the intuitively satisfying interpretation as a distinguishability criterion, i.e., the drift estimate, $\hat{A}$, cannot 'hug' zero forever if we are to distinguish it from zero in finite time. If we adopt an 'energy' viewpoint, the conditions taken together insist that our system have infinite energy in infinite time and finite energy in finite time, respectively. Note that the cost $\{c_t : t \geq 0\}$ satisfies $(A1)$ and $(A2)$ if the conditions $(E1)$ and $(E2)$ are met. If we suppose that the escape conditions are met, and take $I = (a,b)$, with $0 < a < b < 1$, then $\partial_\Pi I = \partial I = \{a, b\}$. Thus, all first exit times based on such continuation intervals will belong to $\mathcal{T}$, i.e., all first exit policies of this type will eventually terminate (P_π-a.s.). Next, consider the following problem,

$$(\mathbb{P}) \quad \begin{array}{lll} (P1): & \frac{1}{2}\pi^2(1-\pi)^2 r_I''(\pi) = -c & \forall \pi \in (0,1); \\ (P2): & r_I(\pi) = e(\pi) & \pi \in \{a,b\}; \\ (P3): & r_I(\pi) < e(\pi) & \pi \notin \{a,b\}, \end{array}$$

with $r_I : [0,1] \to \mathbb{R}$ twice continuously differentiable. A solution to $\mathbb{P}$ therefore consists of two things: an interval $I = (a,b)$ and a function r_I with the above properties. The relationship of $(P2)$ to $(C2)$ and $(P3)$ to $(C3)$ is clear. As for the relationship of $(P1)$ to $(C1)$, note that the Itô formula applied to the function r_I and the Π process yields (P_π-a.s.),

$$r_I(\Pi_t) - r_I(\Pi_0) = \int_0^t r_I'(\Pi_s)\, d\Pi_s + \frac{1}{2}\int_0^t r_I''(\Pi_s)\, d[\Pi,\Pi]_s, \quad (2.13)$$

where the quadratic variation of Π is given by (see 2.12),

$$\begin{aligned} [\Pi,\Pi]_t &= \int_0^t \left(\Pi_s(1-\Pi_s)\hat{A}_s\right)^2 d[\overline{W},\overline{W}]_s \\ &= \int_0^t \Pi_s^2(1-\Pi_s)^2 \hat{A}_s^2\, ds. \end{aligned} \quad (2.14)$$

As a result we can express the second integral in 2.13 as

$$\begin{aligned} \frac{1}{2}\int_0^t r_I''(\Pi_s)\, d[\Pi,\Pi]_s &= \int_0^t \frac{1}{2}\Pi_s^2(1-\Pi_s)^2 r_I''(\Pi_s)\hat{A}_s^2\, ds \\ &= -c\int_0^t \hat{A}_s^2\, ds = -\int_0^t c_s\, ds, \end{aligned} \quad (2.15)$$

using the form of the running cost as chosen in 2.3. Hence 2.13 becomes

$$r_I(\Pi_t) - r_I(\Pi_0) = \int_0^t r_I'(\Pi_s)\, d\Pi_s - \int_0^t c_s\, ds. \quad (2.16)$$

From this last expression we will show that,

$$E_\pi[r_I(\Pi_\tau) - r_I(\Pi_0)] = -E_\pi \int_0^\tau c_s\, ds \quad \forall \tau \in \mathcal{T}, \quad (2.17)$$

which is precisely $(C1)$. Hence, if we can find I and r_I to solve problem $(\mathbb{P})$ then we will have solved the problem posed by $(C1$ - $C3)$ with $r_* = r_I$, and $I_* = I$. To show that 2.17 holds, begin by defining the localizing stopping times

$$\gamma_n := \inf\{t \geq 0 : \Pi_t \notin (\frac{1}{n}, \frac{n-1}{n})\}, \quad (2.18)$$

and the bounding sequence

$$B_n := \sup_{\frac{1}{n} \leq \pi \leq \frac{n-1}{n}} [r_I'(\pi)]^2 \quad \forall n \geq 1. \quad (2.19)$$

Note that $\gamma_n \leq \gamma_{n+1}$, $\gamma_n \to \infty$, and in addition $\gamma_n < +\infty$ P_π-a.s., because Π is assumed to satisfy the escape conditions. Also note that B_n is finite for all $n \geq 1$ because $r_I \in C^2(0,1)$ implies r_I' is locally bounded on $(0,1)$. Since our aim is to show that the first integral on the right hand side of 2.16 is a $(\mathcal{O}_t, P_\pi)$-local martingale. To this end choose any $\tau \in \mathcal{T}$ and compute,

$$E_\pi \int_0^{\tau\wedge\gamma_n} [r_I'(\Pi_s)]^2\, d\langle\Pi,\Pi\rangle_s \leq B_n\, E_\pi\langle\Pi,\Pi\rangle_{\tau\wedge\gamma_n} \quad (2.20)$$

where $\langle\Pi,\Pi\rangle$ is the predictable compensator of Π^2. In view of the fact that Π is a corlol uniformly integrable $(\mathcal{O}_t, P_\pi)$-martingale with $E_\pi\Pi_\infty^2$ finite, it follows directly from the properties of isometric integrals [Wong] that $\langle\Pi,\Pi\rangle$ is also the $(\mathcal{O}_t, P_\pi)$-predictable compensator of $[\Pi,\Pi]$. Hence,

$$\begin{aligned} E_\pi\langle\Pi,\Pi\rangle_{\tau\wedge\gamma_n} &= E_\pi[\Pi,\Pi]_{\tau\wedge\gamma_n} = E_\pi[\Pi_{\tau\wedge\gamma_n}^2 - 2E_\pi\int_0^{\tau\wedge\gamma_n}\Pi_s\, d\Pi_s] \\ &\leq 1, \end{aligned} \quad (2.21)$$

because $\int_0^t \Pi_s\, d\Pi_s$ is clearly an $(\mathcal{O}_t, P_\pi)$-martingale. From this we obtain,

$$E_\pi \int_0^\infty [r_I'(\Pi_s)]^2\, 1\{s \leq \gamma_n\}\, d\langle\Pi,\Pi\rangle_s \leq B_n < +\infty \quad \forall n \geq 1, \quad (2.22)$$

and therefore $\int_0^t r_I'(\Pi_s) 1\{s \leq \gamma_n\}\, d\Pi_s$ is an $(\mathcal{O}_t, P_\pi)$-martingale, i.e., $\int_0^t r_I'(\Pi_s)\, d\Pi_s$ is an $(\mathcal{O}_t, P_\pi)$-local martingale. So from 2.16,

$$\begin{aligned} E_\pi[r(\Pi_{\tau\wedge\gamma_n}) - r(\Pi_0)] &= E_\pi\int_0^{\tau\wedge\gamma_n} r'(\Pi_s)\, d\Pi_s - E_\pi\int_0^{\tau\wedge\gamma_n} c_s\, ds \\ &= -E_\pi\int_0^{\tau\wedge\gamma_n} c_s\, ds \quad \forall n \geq 1. \end{aligned} \quad (2.23)$$

Now, since $\gamma_n \to \infty$, then $\tau\wedge\gamma_n \to \tau$ (P_π-a.s.) because $\tau \in \mathcal{T}$. Also, since Π is a uniformly integrable martingale this implies $\Pi_{\tau\wedge\gamma_n} \to \Pi_\tau$, and thence $r_I(\Pi_{\tau\wedge\gamma_n}) \to r_I(\Pi_\tau)$, from the (assumed) continuity of r_I. In view of $(P2)$ and $(P3)$ we see that r_I is bounded, $(r_I \leq e(\pi_e))$ and so the Bounded Convergence Theorem gives,

$$E_\pi[r_I(\Pi_{\tau\wedge\gamma_n})] \to E_\pi[r_I(\Pi_\tau)]. \quad (2.24)$$

Furthermore, $0 \leq \int_0^{\tau\wedge\gamma_n} c_s\, ds \leq \int_0^{\tau\wedge\gamma_{n+1}} c_s\, ds \to \int_0^\tau c_s\, ds$ (P_π-a.s.) so that the Monotone Convergence Theorem yields,

$$E_\pi\int_0^{\tau\wedge\gamma_n} c_s\, ds \to E_\pi\int_0^\tau c_s\, ds. \quad (2.25)$$

From 2.23, 2.24, and 2.25 we get therefore

$$E_\pi[r_I(\Pi_\tau) - r_I(\Pi_0)] = -E_\pi\int_0^\tau c_s\, ds. \quad (2.26)$$

Since $\tau \in \mathcal{T}$ is arbitrary, 2.17 follows as promised and the connection between $(C1)$ and $(P1)$ is established. As a result, we restate for emphasis that if $(\mathbb{P})$ can be solved then we will have found a pair (r_I, I) which satisfies $(C1$ - $C3)$. Stated another way, if (r_I, I) exists satisfying $(\mathbb{P})$, then $(r_I, I) \equiv (r_*, I_*)$ with equivalence in the sense of essential uniqueness. We now proceed to solve $(\mathbb{P})$. First, fix any $a, b \in \mathbb{R}$ satisfying $0 < a < \pi_e < b < 1$ and define $I := (a,b)$. We do this because it follows necessarily from Lemma 2 that $0 < a_* < \pi_e < b_* < 1$. Next, consider the following ODE,

$$\begin{aligned} \frac{1}{2}\pi^2(1-\pi)^2 r_I''(\pi) &= -c \quad \forall \pi \in (0,1); \\ r_I(a) &= e(a); \\ r_I(b) &= e(b). \end{aligned} \quad (2.27)$$

By elementary ODE theory we see that a unique solution exists; call it $r_I(\pi) = r(\pi; a, b)$ with $r_I \in C^2(0,1)$. Letting $e^+(\pi) := c^0\pi$, $e^-(\pi) := c^1(1-\pi)$, and recalling the definition of the mapping e we have,

$$e(\pi) = \min\{e^+(\pi), e^-(\pi)\}. \quad (2.28)$$

Define the auxiliary function, $g(\pi) := r_I(\pi) - e(\pi)$, and observe

$$g(\pi) = \max\{r_I(\pi) - e^+(\pi), r_I(\pi) - e^-(\pi)\}. \quad (2.29)$$

From 2.27 it is clear that $r_I'' < 0$ and so r_I is (strictly) concave. Hence, g is the maximum of two concave functions, say $g^\pm$ with $g^\pm \equiv r - e^\pm$, and therefore is generically not itself concave (see Fig. 1). Denote by G, G^+, and G^- the hypographs of g, g^+, and g^-, and let $K = Co(G)$, the convex hull of G. Now if $K \backslash G = \phi$, then since $G = G^+ \cup G^-$ either $K \backslash G^+ = \phi$, or $K \backslash G^- = \phi$ or both. From a straightforward case by case analysis this

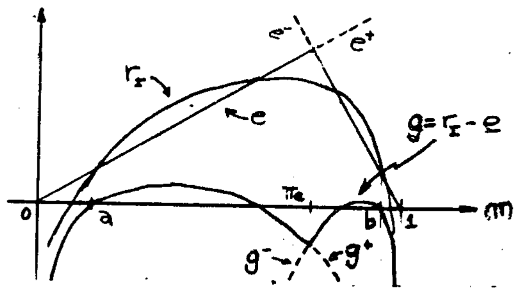

Fig. 1

implies $c^0 c^1 = 0$, which contradicts our assumption that $c^0, c^1 > 0$. Therefore $K \backslash G \neq \phi$ and choosing $z = (\pi_e, g(\pi_e)) \in \partial(K \backslash G)$, we construct the hyperplane at z supported by K. Since $z \in \partial K$, there exist $\lambda \in (0,1)$, $x_* \in \partial G^+$, and $y_* \in \partial G^-$ such that $z = \lambda x_* + (1-\lambda)y_*$. Let $\ell_* : [0,1] \to \mathbb{R}$ denote the line whose graph is this hyperplane (see Fig. 2). We have

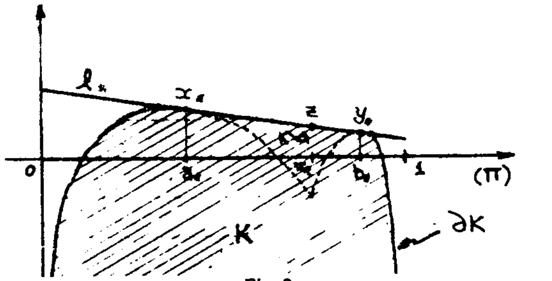

Fig. 2

$x_* = (a_*, \ell_*(a_*)), y_* = (b_*, \ell_*(b_*))$ for some a_*, b_* with $0 < a < a_* < \pi_e < b_* < 1$ so that $g(a_*) = \ell_*(a_*)$ and $g(b_*) = \ell_*(b_*)$. We note that,

$$\ell_*(\pi) \geq g(\pi) = r_I(\pi) - e(\pi) \qquad \forall \pi \in [0,1] \tag{2.30}$$

and if we define, $r_*(\pi) := r_I(\pi) - \ell_*(\pi)$, then obviously,

$$r_*(\pi) \leq e(\pi) \qquad \forall \pi \in [0,1]. \tag{2.31}$$

Moreover,

$$\begin{aligned} r_*(a_*) &= r_I(a_*) - \ell_*(a_*) \\ &= e(a_*) + (r_I(a_*) - e(a_*)) - \ell_*(a_*) \\ &= e(a_*) + g(a_*) - \ell_*(a_*) = e(a_*), \end{aligned} \tag{2.32}$$

and similarly,

$$r_*(b_*) = e(b_*). \tag{2.33}$$

Thus, 2.32 and 2.33 combine to give $(P2)$. Moreover, in view of the strict concavity of r, it follows from the definition of r_* that these are the only two points for which the inequality in 2.31 is not strict, i.e.,

$$r_*(\pi) < e(\pi) \qquad \forall \pi \notin \{a_*, b_*\}, \tag{2.34}$$

which gives $(P3)$. Finally, we point out that,

$$r_*''(\pi) = r_I''(\pi) - \ell_*''(\pi) = -c - 0 = -c, \tag{2.35}$$

which is $(P1)$. Thus, r_* as defined above and $I_* := (a_*, b_*)$ solve $(\mathbb{P})$ and therefore solves the problem posed by $(C1$ -$C3)$. Thus, the Verification Theorem applied to our problem yields the following.

Theorem 3: Assume (E1) and (E2) hold. In the problem of sequential detection based on observations of the process (see 2.1),

$$Y_t = H_\pi \int_0^t A_s\, ds + W_t \qquad t \geq 0, \tag{2.36}$$

with average running cost (see 2.3),

$$E_\pi \int_0^\tau c_s\, ds = E_\pi \int_0^\tau c \hat{A}_s^2\, ds \qquad \forall \tau \in \mathcal{T}, \tag{2.37}$$

and average decision cost $E_\pi[C(H_\pi, \delta)]$ (see 2.5), there exist a_*, b_* unique, with $0 < a_* < \pi_e < b_* < 1$, such that the first exit policy, (τ_*, δ_*), with continuation interval $I_* = (a_*, b_*)$ is Bayesian optimal.

Proof: The existence and uniqueness of a_*, b_* satisfying $0 < a_* < \pi_e < b_* < 1$ has been demonstrated in the preceding discussion, as is the Bayesian optimality of the first exit policy (τ_*, δ_*) given by

$$\tau_* = \inf\{t \geq 0 : \Pi_t \notin I_*\}; \tag{2.38}$$

$$\delta_* = \begin{cases} 1 & \text{if } \Pi_{\tau_*} \geq \pi_e; \\ 0 & \text{if } \Pi_{\tau_*} \leq \pi_e. \end{cases} \tag{2.39}$$

We point out that one can also write,

$$\delta_* = \begin{cases} 1 & \text{if } \Pi_{\tau_*} \geq b_*; \\ 0 & \text{if } \Pi_{\tau_*} \leq a_*, \end{cases} \tag{2.40}$$

since $\pi_e \in I_*$. □

Under the hypothesis that the drift process A is a constant, this theorem has been shown by Shiryayev (1978). However, the demonstration of the existence and uniqueness of a_* and b_* in (Shiryayev, 1978) corresponding to the convexity argument given herein uses complicated analytic methods. In general, there are no closed form solutions for a_* and b_* in terms of c, c^0, and c^1 for this problem, and one must resort to numerical approximation (see below). There is however one special case when closed form solutions are possible, namely when c, c^0, and c^1 are related by, $c^0 = c^1 = (3 - \frac{1}{3} + 2\log 3)\cdot c$, for which a_* and b_* have the simple symmetric solutions, $a_* = \frac{1}{4}$, $b_* = \frac{3}{4}$. Otherwise approximation is necessary, and numerical methods such as the Newton-Raphson technique can be used to find the roots of the two dimensional nonlinear system of equations,

$$\begin{aligned} r_I(a) &= e(a); \\ r_I(b) &= e(b). \end{aligned} \tag{2.41}$$

However, the convexity analysis given above suggests the following simple algorithm to compute a_*, b_* and thence r_*. Choose any a_0, b_0, $0 < a_0 < \pi_e < b_0 < 1$ and solve for r_0 satisfying the ODE 2.27 with $I_0 = (a_0, b_0)$. This is easy to do since one can check that $r(\pi; a, b)$ defined via,

$$r(\pi; a, b) := d(\pi) + [e(a) - d(a)]\frac{b - \pi}{b - a} + [e(b) - d(b)]\frac{\pi - a}{b - a}, \tag{2.42}$$

with,

$$d(\pi) := c(1 - 2\pi)\log\frac{\pi}{1 - \pi}, \tag{2.43}$$

indeed solves 2.27 for any $0 < a < b < 1$ and so we define $r_{I_0}(\pi) := r(\pi; a_0, b_0)$ for all $\pi \in (0,1)$. Next define $g_0 \equiv r_{I_0} - e$ and consider the following recursion, $g_n \equiv g_{n-1} - \ell_n \quad n = 1, 2, \ldots$, where,

$$\ell_n(\pi) := g_{n-1}(a_n)\frac{b_n - \pi}{b_n - a_n} + g_{n-1}(b_n)\frac{\pi - a_n}{b_n - a_n}. \tag{2.44}$$

The iterates a_n, b_n need only satisfy,

$$\begin{aligned} a_n &\in A_n := \{\pi < \pi_e : g_n(\pi) > 0\}; \\ b_n &\in B_n := \{\pi > \pi_e : g_n(\pi) > 0\}. \end{aligned} \tag{2.45}$$

It is simple to check that $A_n \supset A_{n+1}, B_n \supset B_{n+1}$ for all $n \in \mathbb{N}$ and so the algorithm always converges. In fact, if a_n, b_n are chosen in a reasonable manner (e.g., near the true osculation points) then the convergence of $\{a_n\}$ to a_* and $\{b_n\}$ to b_* is quite fast. Our experience indicates about four decimal places of accuracy is achieved in no more than three steps. For completeness we point out that Bayes' optimal cost, ρ, is given by,

$$\rho(\pi) = \begin{cases} r_*(\pi) & \pi \in I_* \\ e(\pi) & \pi \notin I_* \end{cases} \tag{2.46}$$

where $I_* := (a_*, b_*)$, $r_*(\pi) = r_{I_*} := r(\pi; a_*, b_*)$ with r as in 2.42.

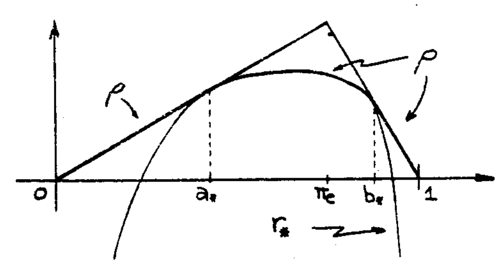

Fig. 3

In Fig. 3 we depict the optimal configuration. Since $r_* \in C^2(0,1)$, it is clear from the figure (or equivalently, from $(P2)$ and $(P3)$) that the derivatives of r_* at a_* and b_* match the derivatives of e at these two points, i.e.,

$$\begin{cases} r_*'(a_*) = e'(a_*); \\ r_*'(b_*) = e'(b_*). \end{cases} \tag{2.47}$$

These are the so-called "smooth pasting" conditions (Shiryayev, 1978) related to the Stephan problem associated with the sequential detection problem of a homogeneous diffusion with average running cost $E_\pi[\tilde{c}\tau], \tilde{c} > 0$. Our formulation includes this case which one can see by choosing

$$A_t = A \neq 0 \qquad \forall t \geq 0, \tag{2.48}$$

i.e., by supposing that Y is a homogeneous diffusion with a constant, nonzero drift coefficient given by A. In this case (E1) and (E2) are trivially satisfied and the average running cost is given by,

$$E_\pi \int_0^\tau c_s ds = E_\pi \int_0^\tau c\hat{A}_s^2 ds = E_\pi[cA^2\tau] = E_\pi[\tilde{c}\tau] \qquad \forall \tau \in \mathcal{T}, \tag{2.49}$$

if we define $\tilde{c} = cA^2 > 0$. Thus, the "smooth pasting" conditions follow necessarily from the conditions $(C1$ - $C3)$. This is significant because the Stephan problems as usually derived in problems of optimal stopping typically have nonunique solutions and one must search for additional conditions which the optimal solution satisfies and which sufficiently constrain the Stephan problem to force its solution to be unique.

Moreover, we have used our approach in solving an optimal stopping problem arising in the sequential detection of conditionally Poisson counting processes (Baras and MacEnany, 1987) where "smooth pasting" does not hold. The same conditions ($C1$ - $C3$) however work just as well for that problem as for the one here. In closing this section we observe that it is no surprise that we obtain the same "smooth pasting" conditions, and indeed the same Bayes' cost, for the apparently more general problem (2.1), as for the homogeneous diffusion case. This is a result of our choice of running cost which as we have pointed out 'collapses down' to the usual cost found in the homogeneous case. Mathematically, this cost simply transforms the homogeneous problem by a random change of time scale. Heuristically, this cost penalizes the detector according to a sample path dependent 'elapsed energy' criterion rather than a uniform 'elapsed time' criterion. In a real-world application one can argue on physical grounds that this choice of running cost is often more reasonable in that it reflects the statistician's modest desire to expect better performance from the detector (faster decisions on average) in 'favorable' problems (large drifts to detect), and to allow worse performance (slower decisions on average) in 'hard' problems (drifts which 'hug' zero). Apparently, such a statistician believes that "not all filtrations are created equal", to put it colloquially. This cost was first used by Shiryayev in the Wald problem of sequential detection in (Liptser and Shiryayev, 1978). (We point out that one can prove the optimality of first exit policies for the Wald case directly from the results contained herein by using the method due to Le Cam (Lehmann, 1958)). The same cost was also used for the Bayesian formulation in (LaVigna, 1986). However, both of these papers employ arguments which work in only the diffusion case whereas as our approach works in both the diffusion and jump process cases. We note in passing that the (binary) sequential detection problem is in fact a filtering problem for a partially observed (two-state) Markov Chain, $H_{\pi,t}$, with initial distribution (in our notation),

$$P_\pi\{H_{\pi,0} = 1\} = \pi = 1 - P_\pi\{H_{\pi,0} = 0\}, \tag{2.50}$$

which is degenerate in that both states are absorbing, i.e., the chain makes no transitions. With this understanding, Theorem 1 can be extended to include the more general situation. We have applied such an extension to the disruption problem (MacEnany, 1987), in which $\pi = 0$ and the underlying Markov Chain makes a single transition, and have obtained the optimal solution. Again, there was no need to search for additional conditions as would be necessary in the strict Stephan problem approach (Shiryayev, 1978).

CONCLUSION

We present a verification-type theorem for the Bayesian sequential detection problem which asserts that if there exists an interval and a function satisfying the three conditions of the theorem, then this interval defines the optimal first exit policy and the function defines the Bayes' optimal cost. We apply this result to the Bayesian sequential detection problem for a generalized diffusion and show how the three conditions lead directly to a boundary-value problem of Stephan-type which has a unique solution pair. The proof of the existence and uniqueness of this solution is based on geometric notions of convex sets thereby avoiding complicated analytical arguments. We also provide an algorithm derived from these notions which converges quickly to the optimal solution. We indicate that our methods extend to jump-type processes and problems of disruption in which the analytical approach is excessively complex or hopeless. These applications will appear elsewhere.

REFERENCES

Baras, J.S., and D.C. MacEnany (1987). Sequential Testing on the Rate of a Counting Process. Proc 26th IEEE Conf. on Dec. and Contr. 1592-1593

LaVigna, A. (1986). Real Time Sequential Detection for Diffusion Signals. Master's Thesis, University of Maryland.

Lehmann, E.L. (1958). *Testing Statistical Hypotheses.* John Wiley and Sons, New York.

Lipster, R.S., and A.N. Shiryayev (1978). *Statistics of Random Processes II: Applications.* Springer-Verlag, New York-Heidelberg-Berlin-Toyko.

MacEnany, D.C. (1987). Bayesian Sequential Hypothesis Testing. Tech. Rep. SRC-TR87-107, Systems Research Center, Univ. of Maryland

Shiryayev, A.N. (1977), *Optimal Stopping Rules.* Springer-Verlag, New York-Heidelberg-Berlin-Toyko.

Wong, E., and B. Hajek (1985). *Stochastic Processes in Engineering Systems.* Springer-Verlag, New York-Heidelberg-Berlin.

APPENDIX

Using the notation of the first section, let P_0^t, P_1^t, P_π^t denote the respective $\mathcal{O}_t$-restrictions of the measures P_0, P_1, P_π and observe that $P_1^t \ll P_\pi^t$ and hence,

$$P_1\{O\} = \int_O \frac{dP_1^t}{dP_\pi^t} dP_\pi \qquad \forall O \in \mathcal{O}_t, \tag{A.1}$$

where $\frac{dP_1^t}{dP_\pi^t}$ denotes the Radon-Nikodym derivative. We have the following lemma.

Lemma. For all $t \geq 0$, $\Pi_t = \pi \frac{dP_1^t}{dP_\pi^t}$, except on $\mathcal{O}_t$-sets of P_π-measure zero.

Proof: Let $O \in \mathcal{O}_t$ and compute,

$$\int_O \pi_t \, dP_\pi = \int_O E_\pi[H_\pi|\mathcal{O}_t] \, dP_\pi = \int_O H_\pi \, dP_\pi = P_\pi\{\{H_\pi = 1\} \cap O\}$$

$$= \pi P_1\{O\} = \int_O \pi \frac{dP_1^t}{dP_\pi^t} \, dP_\pi, \tag{A.2}$$

using A.1. Since $O \in \mathcal{O}_t$ is arbitrary we are done. □

We use this to prove the following.

Proposition. Let γ be some $\mathcal{O}$-measurable random variable. Then for all $t \geq 0$,

$$E_\pi[H_\pi \gamma|\mathcal{O}_t] = \Pi_t \, E_1[\gamma|\mathcal{O}_t],$$

except on $\mathcal{O}_t$-sets of P_π-measure zero.

Proof: Let $O \in \mathcal{O}_t$ and compute,

$$\begin{aligned} \int_O E_\pi[H_\pi \gamma|\mathcal{O}_t] \, dP_\pi &= \int_O H_\pi \gamma \, dP_\pi \\ &= \int_{O \cap \{H_\pi = 1\}} \gamma \, dP_\pi \\ &= \int_O \gamma \pi \, dP_1 \\ &= \int_O E_1[\gamma|\mathcal{O}_t] \pi \, dP_1 \\ &= \int_O E_1[\gamma|\mathcal{O}_t] \pi \frac{dP_1^t}{dP_\pi^t} dP_\pi \\ &= \int_O E_1[\gamma|\mathcal{O}_t] \Pi_t \, dP_\pi, \end{aligned} \tag{A.3}$$

using the lemma. Again, the arbitrariness of $O \in \mathcal{O}_t$ gives us the result. □

We conclude this appendix by connecting the Π process to the likelihood ratio process Λ given by,

$$\Lambda_t := \frac{dP_1^t}{dP_0^t} = E_0\left[\frac{dP_1}{dP_0}|\mathcal{O}_t\right] \qquad t \geq 0, \tag{A.4}$$

assuming of course that $P_1 \ll P_0$.

Theorem. Suppose $P_1 \ll P_0$. Then for all $t \geq 0$,

$$\Pi_t = \frac{\pi \Lambda_t}{1 - \pi + \pi \Lambda_t},$$

except on $\mathcal{O}_t$-sets of P_π-measure zero.

Proof: From the lemma,

$$\Pi_t = \pi \frac{dP_1^t}{dP_\pi^t} \qquad \forall t \geq 0, \tag{A.5}$$

except on $\mathcal{O}_t$-sets of P_π-measure zero. We observe that $P_1^t \ll P_0^t$. Using the Radon-Nikodym theorem successively one obtains the following sequence of steps,

$$\begin{aligned} \Pi_t = \frac{\pi \, dP_1^t}{\pi P_1^t + (1-\pi) dP_0^t} &= \frac{\pi \frac{dP_1^t}{dP_0^t}}{1 - \pi + \pi \frac{dP_1^t}{dP_0^t}} \\ &= \frac{\pi \Lambda_t}{1 - \pi + \pi \Lambda_t}, \end{aligned} \tag{A.6}$$

which is the result. □

From A.6 it is clear that,

$$\Pi \equiv \begin{cases} 0 & \text{if } \pi = 0; \\ 1 & \text{if } \pi = 1. \end{cases} \tag{A.7}$$

If $\pi \in (0, 1)$ therefore, we can invert A.5 and obtain,

$$\Lambda_t = \frac{1 - \pi}{\pi} \left(\frac{\Pi_t}{1 - \Pi_t}\right) \qquad \forall t \geq 0. \tag{A.8}$$

In view of this relationship it is clear that first exit policies based on Π or Λ are equivalent in terms of optimality criteria, only their continuation intervals differ.